arXiv:1607.02964v1 [astro-ph.CO] 11 Jul 2016

Agnostic cosmology in the CAMEL framework

S. Henrot-Versillé, O. Perdereau, S. Plaszczynski *, B. Rouillé d'Orfeuil, M. Spinelli and M. Tristram.

Laboratoire de l'Accélérateur Linéaire, Univ. Paris-Sud, CNRS/IN2P3, Université Paris-Saclay, Orsay, France.

July 12, 2016

Abstract

Cosmological parameter estimation is traditionally performed in the Bayesian context. By adopting an "agnostic" statistical point of view, we show the interest of confronting the Bayesian results to a frequentist approach based on profile-likelihoods.

To this purpose, we have developed the *Cosmological Analysis with a Minuit Exploration of the Likelihood* (CAMEL) software. Written from scratch in pure C++, emphasis was put in building a clean and carefully-designed project where new data and/or cosmological computations can be easily included.

CAMEL incorporates the latest cosmological likelihoods and gives access *from the very same input file* to several estimation methods:

- A high quality Maximum Likelihood Estimate (a.k.a "best fit") using MINUIT,
- profile likelihoods,
- a new implementation of an Adaptive Metropolis MCMC algorithm that relieves the burden of reconstructing the proposal distribution.

We present here those various statistical techniques and roll out a full use-case that can then used as a tutorial. We revisit the ΛCDM parameters determination with the latest *Planck* data and give results with both methodologies. Furthermore, by comparing the Bayesian and frequentist approaches, we discuss a "likelihood volume effect" that affects the optical reionization depth when analyzing the high multipoles part of the *Planck* data.

The software, used in several *Planck* data analyzes, is available from http://camel.in2p3.fr. Using it does not require advanced C++ skills.

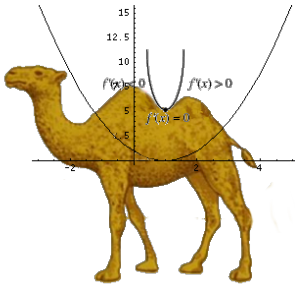

*corresponding author plaszczy@lal.in2p3.fr

Contents

1 Introduction **3**

2 The cosmological parameters estimation (CPE) pipeline **3**

3 A full `CAMEL` use-case **5**
- 3.1 Building and validating the parameter file . . . 5
  - 3.1.1 Parameter file . . . 5
  - 3.1.2 Adapting the Boltzmann solver precision . . . 7
  - 3.1.3 Estimating the wall time . . . 7
- 3.2 The Maximum Likelihood Estimator (MLE), aka ”best fit” . . . 8
  - 3.2.1 Minimisation . . . 8
  - 3.2.2 Trying different starting points . . . 10
  - 3.2.3 The Hessian matrix . . . 11
- 3.3 Profile-likelihoods . . . 11
  - 3.3.1 Properties . . . 11
  - 3.3.2 Profile-likelihoods in practice . . . 12
  - 3.3.3 Profile-likelihood summary . . . 14
  - 3.3.4 Results on ΛCDM cosmological parameter . . . 15
- 3.4 MCMC . . . 15
  - 3.4.1 Real-life Metropolis . . . 15
  - 3.4.2 Adaptive Metropolis (AM) . . . 17
  - 3.4.3 Runing AM . . . 18
  - 3.4.4 Convergence . . . 19
  - 3.4.5 MCMC analysis . . . 21

4 Comparison of the results **23**
- 4.1 ΛCDM parameters . . . 23
- 4.2 Likelihood volume effects . . . 24

5 Summary **26**

1 Introduction

Since the 2000's (Christensen et al. 2001), the adoption of Monte Carlo Markov Chain (MCMC) techniques by the cosmological community has promoted the acceptance of Bayesian methodology for parameter estimation. The possible (although inefficient) sampling of a high dimensional space allows, invoking Bayes theorem, to reconstruct posterior distributions assuming some prior degree of belief (rarely discussed). MCMC usage was popularized in cosmology by the `CosmoMC` package (Lewis & Bridle 2002) which uses an optimized version of the Metroplis-Hastings algorithm (see also `MontePython` (Audren et al. 2013)). There exist today several implementations of other algorithms, as `MultiNest` (Feroz et al. 2009) or `PolyChord` (Handley et al. 2015) for Nested Sampling (Skilling 2004), or `emcee` (Foreman-Mackey et al. 2013) for the affine invariant ensemble sampler (Goodman & Weare 2010). Several of them are packaged within the `CosmosSIS` package (Zuntz et al. 2015) and some of them were compared in Allison & Dunkley (2014). They are essentially written in `Fortran/python` and/or wrapped into multiple languages. All of them adopt the Bayesian paradigm (see Hobson et al. (2010) for a review).

In other communities, as High Energy Physics, the frequentist approach is more traditional. The reason for its absence in cosmology maybe lies in the difficulty of building precise profile-likelihoods because of the numerical noise inherent to the Boltzmann solver computations (Sect.3.1.2). It was however shown to be feasible on real *Planck* data using an accurate minimizer and tuning the precision of the Boltzmann computations (Planck Collaboration Int. XVI 2014).

We wish to avoid the ideological debate about "who is right", and focus on the interest of confronting both methodologies. It is the goal of the `CAMEL` package to provide tools to compare the Bayesian and frequentist analyzes in order to better understand for instance the Bayesian priors effects or the shape of the likelihood. We choose to write it in pure `C++` since we consider it as a proper language to develop a robust, clean and long term project. By `C++` we mean an object-oriented code with some level of abstraction and well defined design patterns in order to build a clean interface that anticipates future developments. Fortunately users do not need to know what is under the hood and one can plug in easily any new Boltzmann solver and/or likelihood.

Certainly the best way to present `CAMEL` is by working out a full use-case. We will revisit the ΛCDM cosmological parameters estimation with *Planck* data using the `Hillipop` likelihood (Planck Collaboration XI 2015) and compare the Bayesian and frequentist results. Not only will we show how to obtain precise best-fits and profile-likelihoods, but also propose a new implementation of an Adaptive scheme for the Metropolis MCMC algorithm that relieves the pain of reconstructing the proposal. Sect. 2 is a quick overview of the building blocks of parameter estimation. Then Sect. 3 explains in minute-details how to obtain best-fits (Sect.3.2), profile-likelihoods (Sect.3.3) and produce MCMC with the new Adaptive algorithm (Sect. 3.4). For each method, the basics will be reviewed. This part concretely shows how to produce all the results and can therefore be used as a tutorial. Sect 4 finally shows the interest of comparing both methods focusing on the study of the likelihood "volume effects". These results, partially available in Planck Collaboration Int. XVI (2014) and Planck Collaboration XIII (2015), were however never compiled and detailed. The results on the volume effects that affect *Planck*'s high-ℓ likelihoods are new and complement Couchot et al. (2015).

2 The cosmological parameters estimation (CPE) pipeline

A CPE analysis can be factorized into 3 distinct pieces (see Fig. 1):

1. A code to compute some observables (as a set of C_ℓ values for the CMB or $P(k, z)$ for the matter power spectra) given a set of cosmological parameters (called Ω) within a theoretical framework: $X(\Omega)$. In cosmology, the numerical codes that solve the background+perturbation equations are often called "Boltzman solvers". In `CAMEL` such a tool is called an **Engine**. Today, `class` (`http://class-code.net/`) and `camb` (`http://camb.info`) are widely used but any analytic parametrization of some cosmological observable enters this category too.

2. The experimental data, which in their most complete statistical form are given in the form of a **likelihood**, i.e a function of X + generally some extra nuisance parameters (called ν). There is an implicit conversion here to cosmological parameters through the Engine: $\mathscr{L}(\Omega, \nu) = \mathscr{L}(X(\Omega), \nu)$. In the following we will also often use the χ^2 function which is by definition $\chi^2 \equiv -2 \ln \mathscr{L}$.

3. One then wants to extract some information about the parameters themselves (including the nuisance ones), so you need some **statistical methodology**. In the realm of *estimators*, the Maximum Likelihood Estimator (MLE) is the King since it enjoys several nice properties (as being unbiased and of minimal variance asymptotically). It is also referred to as a "best fit" meaning the set of $(\hat{\Omega}, \hat{\nu})$ values that gives the maximal likelihood value (or minimum χ^2 one). This is a *point estimate* and is a classical statistical notion. Since it is not very practical to realize what happens in a multi-dimensional space, we focus on individual (or a pair of) parameters : this is the domain of *interval estimation*. Bayesian's will extract the p.d.f of the true parameters (posterior distribution), while frequentists will make statements on data covering the true (but unknown) values through the use of profile-likelihoods.

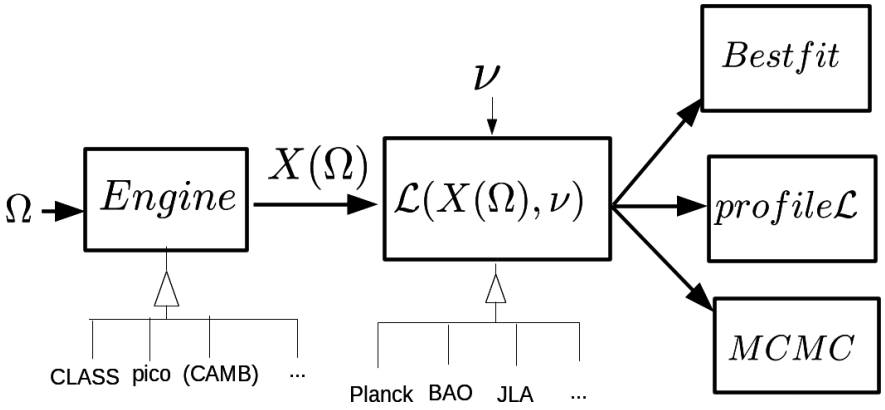

Figure 1: Schematics of the CPE pipeline. From left to right: Ω denotes the cosmological parameters that feed the Engine which computes some observable(s): $X(\Omega)$. ν denotes the nuisance parameters of the likelihood (or any external data necessary to its computation). The last 3 vertical boxes show different statistical methods. Some implementations of engines and likelihoods available in `CAMEL` are derived from the base abstract classes.

`CAMEL` implements these 3 aspects through the abstraction mechanism allowed by Object-Oriented languages. Currently the following concrete implementations (all in `C/C++`, including the likelihoods) are provided:

1. Engines:
   - `class` (Blas et al. 2011)
   - `pico` (Fendt & Wandelt 2007)
2. Likelihoods:
   - CMB: *Planck* (Planck Collaboration XI 2015) *Hillipop* (Tristram et al. 2005; Planck Collaboration XI 2015), *SPT* (Story et al. 2013; Reichardt et al. 2012; George et al. 2015), *ACT* Das et al. (2014),
   - CMB-lensing: *Planck* (Planck Collaboration XV 2016),
   - BAO: *6dF* (Beutler et al. 2011), *WiggleZ* (Blake et al. 2011), *SDSS* (Padmanabhan et al. 2012; Anderson et al. 2012, 2014): 1D and anisotropic,
   - SN1a : *JLA* (Betoule et al. 2014).
3. Statistics: best-fit, profile-likelihood, MCMC.

See the Appendix for more details.

3 A full CAMEL use-case

3.1 Building and validating the parameter file

Estimating cosmological parameter (with a full Boltzman solver) is a CPU demanding task and can hardly be performed on your laptop: one needs to rely on a computer cluster. But before sending a lot of jobs to your farm, you should *define* and *test* your **parameter file** which describes all the cosmological and nuisance parameters (varied or fixed) , data , Engine and statistical method. All the tests -in this section- can be run interactively. We choose as a benchmark the determination of ΛCDM parameters *Planck* high+low ℓ likelihoods using the `class` engine.

3.1.1 Parameter file

We write in this part a parameter file named `hlpTT_bflike_LCDM.par`.

First comes the choice of the engine:

```
engine=class
```

Listing 1: Definition of the engine

Then comes the choice of the likelihoods. Here we will use the *Planck* `Hillipop` likelihood (for the high-ℓ part) and `bflike` for the low-ℓ one (both discussed in Planck Collaboration XIII (2015)).

```
###############################################################################
#hiell
HiLLiPOP=HiLLiPOP/DX11dHM_superExt_CO_TT.lik
#lowl
clikfile=planck_data/low_l/bflike/lowl_SMW_70_dx11d_2014_10_03_v5c_Ap.clik
###############################################################################
```

Listing 2: Defining the likelihoods

Then we define the parameters: the cosmological ones (i.e those entering the engine) must be stamped with the "cosm" keyword and follow the syntax of the Engine(here we follow the `class` terminology). Nuisances, that depends on the likelihoods, should use the "nui" keyword. Each parameter that will be varied ("par") or not ("fix") should be given an initial value (no necessarily very precise), an expected standard deviation and some boundaries for finite exploration (which can be viewed as the flat prior borders for MCMC sampling or simply a box within which the MLE is searched).

```
##############PARAMETERS#################################################
#cosmological parameters
#par/fix  name              type      init          sigma      min      max
par    omega_b             cosm      0.02224       0.00027    0.017    0.027
par    omega_cdm           cosm      0.1192        0.0026     0.09     0.15
par    100*theta_s         cosm      1.0418        0.6E-04    1.03     1.05
par    tau_reio            cosm      0.07          0.13E-01   0.01     0.20
par    log(10^10A_s)       cosm      3.07          0.025      2.7      3.5
par    n_s                 cosm      0.96          0.0070     0.9      1.1

#nuisances for hilipop
par    A_planck            nui       1               0.001      0.9     1.1
par    c0                  nui       0.              0.001     -0.05    0.05
par    c1                  nui       0.              0.001     -0.05    0.05
fix    c2                  nui       0.              0.001     -0.05    0.05
par    c3                  nui       0.              0.001     -0.05    0.05
par    c4                  nui       0.004           0.001     -0.05    0.05
par    c5                  nui       0.004           0.001     -0.05    0.05
par    Aps100x100          nui       2.5E-04         1E-05      0.0     0.1
par    Aps100x143          nui       1.1E-04         7E-06      0.0     0.1
par    Aps100x217          nui       9.9E-05         6E-06      0.0     0.1
par    Aps143x143          nui       4.7E-05         2E-06      0.0     0.1
par    Aps143x217          nui       3.1E-05         3E-06      0.0     0.1
```

```
par    Aps217x217            nui      7.6E-05       6E-06      0.0      0.1
par    Asz                   nui      1              0.1       0.0      10
par    Acib                  nui      1.             0.1       0.0      10
par    AdustTT               nui      1               0.1      0.0      2
fix    AdustPP               nui      0.00          0.1       0.0      2
fix    AdustTP               nui      0.00          0.1       0.0      2
par    Aksz                  nui      0.00          1.0       0.0      10
par    Aszxcib               nui      0.00          1.0       0.0      10
#########################################################################
```

Listing 3: Parameters definition for the `Hillipop+bflike` likelihoods

Next you may add some external constraints (or "priors" in Bayesian vocabulary). The most common one is the 1D Gaussian, which is illustrated here: for instance we add a Gaussain likelihood to "AdutTT" centered on 0 with $\sigma = 0.2$.

```
#priors#################################
gauss1     AdustTT   1  0.2
gauss1     Acib      1  0.2
gauss1     A_planck  1  0.0025

gauss1     c0         0      2e-3
gauss1     c1         0        2e-3
gauss1     c3         0        2e-3
gauss1     c4         0.0025       2e-3
gauss1     c5         0.0025       2e-3
#########################################
```

Listing 4: Adding Gaussian priors

Then you may need to tune the details of the Boltzman solver. For `class` you can do it with the keyword `class` which passes the following argument as a strings to the solver. For instance to define the non-cold-dark matter setup to a degenerate massive neutrino with 0.06 eV (the upper limit in normal hierarchy measured from oscillation experiments) one may use:

```
#class setup############################
class      N_ncdm                  1
class      m_ncdm                  0.06
class      N_eff                   2.046
########################################
```

Listing 5: Defining extra `class` setup

Note that in `class`, $N_{\rm eff}$ corresponds only to the pure radiation today which is why it is not the famous 3.046 value). One may also pass these parameters as previously with the "fix" keyword, but this better decouples the variables from their environment (and avoids repeating fixed numbers in the output files). There are some cases where strings to `class` is mandatory:

```
#################################################
class      k_pivot                 0.05
class      lensing                 yes
class      sBBN\ file              bbn/sBBN.dat
#################################################
```

Listing 6: Direct transmission to `class` of string parameters

(Note the \ in the last line to indicate a trailing space)

Finally we define the file that contains some high-precision settings when we wish to compute precisely the MLE or profile with `class`. Indeed as we will see later getting a precise χ^2 minimum is very challenging with a method that cannot rely on analytical gradients. More exactly Boltzmann solvers do not lead to a likelihood that is continuous in the parameters space because they are the result of many numerical computations with some limited accuracy (as cutoffs in the integrals). `class`, beside being written in clean C, is interesting in this respect since it allows to specify

in a single place all the precision parameters that are used. We provide a set of pre-defined precision files (which result from a trade-off between accuracy and CPU-time and are presented in Planck Collaboration Int. XVI (2014)) and suggest to use the following one:

```
########################################
precisionFile=class_pre/hpjul2.pre
########################################
```

Listing 7: Using a higher precision setup for `class` computations

when performing MLE or profile-likelihood studies. Default settings can be used when doing MCMC, because the algorithm smooths the details of the likelihood function. More in the next section.

The parameter file is ready and we call it `hlpTT_bflike_LCDM.par`. To test it you can run

```
writeChi2 hlpTT_bflike_LCDM.par
```

which parses your file and just writes the χ^2 value corresponding the the given parameter value (no fit there). If the program goes to the end, you most probably have some valid file.

3.1.2 Adapting the Boltzmann solver precision

While in practice using the `hpjul2.pre` is generally sufficient, we show here how you can get an idea of the numerical noise in your analysis. We recall that the difficulty for finding a really precise multi-dimensional χ^2 minimum comes from the fact that the function $\frac{\partial \mathscr{L}}{\partial \Omega_i}$ is not continuous because we use a full Boltzman solver to compute the observable X in $\mathscr{L}(X(\Omega))$. We can get some idea of the level of numerical noise, by scanning linearly one variable (with all other ones fixed) and studying the smoothness of the $\chi^2(\theta_i)$ values. This is the goal of the `ScanParam` program which fixes all but one parameters from the file, vary linearly the free one and records the χ^2 values in a text file. Let's see how this works for instance on the parameter ω_b. Run:

```
ScanParam hlpTT_bflike_LCDM.par omega_b 0.020 0.025 50
```

where the last 3 values indicates how the variable is scanned (here from 0.02 to 0.025 with 50 equidistant points). The upper plot of Fig. 2 shows the χ^2 values, in blue when using the default `class` precision parameters (ie. no particular precision file) and in red with the `hpjul2.pre` one. Then the black curve shows a smooth interpolation between the points, and the bottom plot the residual of each point w.r.t to it. We then see that the numerical noise with the `hpjul2.pre` file (in red) is much lower than for the default case (in blue). Using such a precision allowed in all cases we studied to determine an accurate χ^2 minimum.

3.1.3 Estimating the wall time

It can be useful before running jobs to estimate the time taken by each iteration. By iteration we mean the complete processing of one model. For this you can run interactively a few iterations of the `Minimize` executable adding some verbosity to your parameter file with:

```
verbose=true
```

This is done with

```
Minimize myparfile output
```

where "output" is any name.

It will print among many other information the time used for each iteration (at the end of each line). The CPU time for `hlpTT_bflike_LCDM.par` typically varied between 5s on 8 cores to 2s on 24 ones (with the `hpjul2` precision file) and is only limited by the scalability of the Boltzman solver.

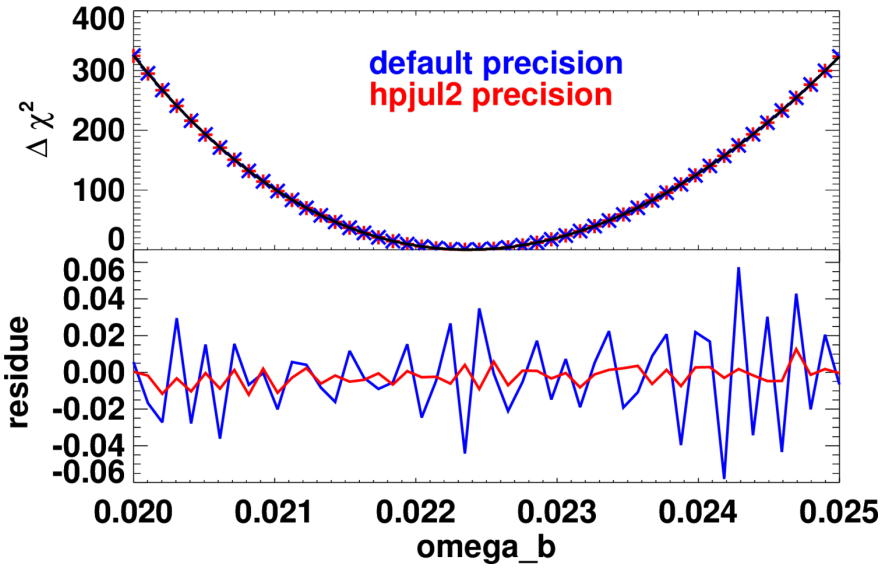

Figure 2: The upper plot shows the χ^2 values obtained from the `ScanParam` program using the `hlpTT_bflike_LCDM.par` setup, in blue when not specifying any precision file for `class` and in red using the hpjul2 one. A smooth fit is superimposed in black. The lower plot shows the residuals in both configurations and reflects the numerical noise of the Boltzmann Engine.

3.2 The Maximum Likelihood Estimator (MLE), aka "best fit"

The MLE enjoys many nice properties that can be reviewed e.g in James (2007). One that is not always well known (and that also applies to the likelihood ratio statistics) is the MLE invariance which states that if $\hat{\theta}$ is the MLE of θ , then the MLE of any function of θ, $y = f(\theta)$ is $\hat{y} = f(\hat{\theta})$.

For instance in the ΛCDM model, there are 6 free cosmological parameters often chosen as $\Omega = \omega_b, \omega_c, \theta, n_s$ $\ln(10^{10}A_s), \tau$ (see `hlpTT_bflike_LCDM.par`). If instead of the angular size of the sound horizon (θ) one uses the Hubble constant (H_0), the minimization leads exactly to the same values for all the other parameters. Furthermore, if the minimization was performed on the Ω set, one does not need to redo it with H_0. From the invariance property you just need to run the Boltzmann solver on the $\hat{\Omega}$ solution to obtain the $\hat{H}_0$ best fit solution.

3.2.1 Minimisation

The goal is to measure in the parameter space where the likelihood reaches its maximum value or equivalently the χ^2 its minimum (which is preferred for numerical reasons). As discussed in Sect.3.1.2 obtaining a precise multi-dimensional minimum (here on 23 parameters but it can be more) is challenging since

- we don't have analytical first order derivatives,
- Boltzmann engine computations brings up numerical noise.

Many softwares exists for multi-dimensional optimization. We chose to work with the `MINUIT`[1] package, a renowned minimizer often used in High Energy Physics. Starting in the 70's (James & Roos 1975), it was permanently improved and rewritten into C++. This version, named `MINUIT2` is embedded within `CAMEL`. `MINUIT` is a large toolbox, including many features (as releasing or constraining parameters on the fly), and working with it requires some degree of experimentation. We tuned it for cosmology with the `MIGRAD` algorithm that is a "variable metric

[1] http://seal.web.cern.ch/seal/work-packages/mathlibs/minuit/index.html

method" (one changes the metric of the system not the data, see eg. Davidon & Laboratory (1959)). In most cases you won't need to change the default setting.

To run the minimization as a batch job, one uses the `Minimise.sh` script (note the ".sh"): [2]

```
Minimize.sh /path/to/hlpTT_bflike_LCDM.par
```

This

1. checks your configuration,
2. creates a directory named hlpTT_bflike_LCDM_min where you launched the script, which will contain the outputs,
3. sends the (OpenMP) job.

You then wait for the job completion... In our case, the minimization converged in about 3000 iterations, which, given the wall times measured in Sect.3.1.3 correspond typically to 1.5h (on 24 cores) or 4h on 8. This time is only limited by `class` computations not the minimization algorithm.

One then obtain the job log-file and the `bestfit.dat` and `covmat.data` files. You can examine the log-file and not worry too much about `class` error messages (if they are sparse!). They are often due to the lack of convergence of some numerical algorithms for some "extreme" combination of parameters: `CAMEL` catches these errors and the minimizer moves away from those regions. You should also not worry about such a message:

```
WARNING: Minuit did not converge.
```

which is due, for `class`, to the numerical discontinuities discussed in Sect 3.1.2. Then near the end you will find the interesting results.

```
# ext. ||   Name    ||   type  ||     Value     ||  Error +/-

0   ||   omega_b   ||  limited ||0.02221669365664     ||0.0001757676645765
1   ||   omega_cdm ||  limited ||0.1193778692005      ||0.001014938068844
2   ||   100*theta_s||  limited ||1.041727722902       ||0.000349347668509
3   ||   tau_reio  ||  limited ||0.07124770721867     ||0.002294613708248
4   ||   log(10^10A_s)||  limited ||3.068270305092       ||0.00899043822157
5   ||   n_s       ||  limited ||0.9657674074055      ||0.006549911640786
6   ||   A_planck  ||  limited ||1.000680632922       ||0.002480642264396
7   ||   c0        ||  limited ||0.001487875696301    ||0.001638396021085
8   ||   c1        ||  limited ||0.001440151762443    ||0.00163998060017
9   ||   c2        ||   const  ||0                    ||
10  ||   c3        ||  limited ||-0.0007967689027367  ||0.00147758139892
11  ||   c4        ||  limited ||0.001778114967395    ||0.001762250083468
12  ||   c5        ||  limited ||0.001921014220322    ||0.001759624112839
13  ||   Aps100x100||  limited ||0.000255535488113    ||3.31867430043e-05
14  ||   Aps100x143||  limited ||0.0001108863113654   ||2.529138206055e-05
15  ||   Aps100x217||  limited ||7.71374529739e-05    ||1.641760577019e-05
16  ||   Aps143x143||  limited ||4.429855974166e-05   ||2.030311938827e-05
17  ||   Aps143x217||  limited ||3.222589327638e-05   ||1.452202004259e-05
18  ||   Aps217x217||  limited ||7.645366758567e-05   ||1.374833402066e-05
19  ||   Asz       ||  limited ||1.322514028996       ||0.5973930819106
20  ||   Acib      ||  limited ||0.9397363536028      ||0.1319397001748
21  ||   AdustTT   ||  limited ||0.9832158679645      ||0.09576457283827
```

[2] Scripts are located under the `work/batch` directory and we propose some implementations for cc-inp3 (`cc.in2p3.fr`) and NERSC (`nersc.gov`); they should be simple to adapt to other sites.

```
22  ||   AdustPP   ||   const  ||0                    ||
23  ||   AdustTP   ||   const  ||0                    ||
24  ||   Aksz      ||  limited ||1.634684530752       ||5.729577959257
25  ||   Aszxcib   ||  limited ||8.295195641495e-08   ||8.899445674664
```

which gives the values of the best fit and some (crude) estimate of the "error bars" from the Hessian matrix (that should not be trusted too much, more in Sect. 3.2.3). Also of uttermost importance is the value of the $\chi^2_{\rm min}$ at this location which appears as

```
minimum function Value: 20448.12723918
```

and at the very end of the file you obtain the breakdown of the individual $\chi^2_{\rm min}$ values and the total wall-time. Here we obtain:

```
===== lowl_SMW_70_dx11d_2014_10_03_v5c_Ap.clik: chi2=10495
===== HiLLiPOP:DX11dHM_superExt_CO_TT.lal: chi2=9951.47
Chi2 from HiLLiPOP: 9951.47 ndof: 9556
===== CMB_all: chi2=20446.5
===== Gauss1 constraint on AdustTT(1+-0.2): chi2=0.00704268
===== Gauss1 constraint on Acib(1+-0.2): chi2=0.0907927
===== Gauss1 constraint on A_planck(1+-0.0025): chi2=0.0741218
===== Gauss1 constraint on c0(0+-0.002): chi2=0.553444
===== Gauss1 constraint on c1(0+-0.002): chi2=0.518509
===== Gauss1 constraint on c3(0+-0.002): chi2=0.15871
===== Gauss1 constraint on c4(0.0025+-0.002): chi2=0.13028
===== Gauss1 constraint on c5(0.0025+-0.002): chi2=0.0838061
TIMER TOTAL TIME=16268.6 s      =271.144 min    =4.51907h
===== Chi2Combiner: chi2=20448.1
```

Now for simplicity we included the `bestfit.awk` script [3] which parses the whole output file and prints out the interesting things. Use it as:

```
awk -f tools/awk/bestfit.awk logfile
```

The files `bestfit.dat` and `covmat.dat` are also interesting excerpts of the log-file that can be easily read with any interactive analysis software.

3.2.2 Trying different starting points

It can be wise to test whether one could not get a slightly better (lower) $\chi^2_{\rm min}$ value through a different path (recall the numerical Engine noise may perturb the gradient descent). To this end, we would like to shoot different starting points and see what we obtain. This can be done specifying a job "range" in the form "I-J". For example:

```
Minimize.sh path/to/hlpTT_bflike_LCDM.par 1-10
```

creates the `hlpTT_bflike_LCDM_bf/` directory (that contains everything) and sends 10 jobs to the cluster through what is referred to as an "array-job".

Each job is independent, so can be run simultaneously with all the others finally increasing little the user wall-time. Job number 1 is the standard minimization described previously. Others will perform some uniform random shuffle of each parameter within the [min,max] bounds of the parameter file. These shots are reproducible in the sense that another "1-10" run would yield the same configuration. Use "11-20" if you want 10 more trials...

Now we have the outputs of jobs 1 to 10 in the same directory. Which is the best? the one with the lowest $\chi^2_{\rm min}$. We provide a script in order to help finding it. Go into `hlpTT_bflike_LCDM_bf/` and run

[3] analysis scripts are located in the `work/tools` directory

```
awk -f tools/awk/zebest.awk *.o*
```

where "*.o*" is a pattern of all the log-files files (and may be different at your site). Here is an example of what you can obtain (given the random seed do not expect exactly the same)

```
hlpTT_bflike_LCDM_bf.o26551546.1 chi2=20448.127239 it=2734 t=5.950475s tot=4.519056h  errors=2
hlpTT_bflike_LCDM_bf.o26551546.10 chi2=20448.205740 it=4706 t=6.512537s tot=8.513333h  errors=13
hlpTT_bflike_LCDM_bf.o26551546.2 chi2=20448.276149 it=3191 t=5.881981s tot=5.213722h  errors=15
hlpTT_bflike_LCDM_bf.o26551546.3 chi2=20448.148907 it=5165 t=6.357231s tot=9.120861h  errors=20
hlpTT_bflike_LCDM_bf.o26551546.4 chi2=20448.199447 it=3300 t=6.380727s tot=5.849000h  errors=56
hlpTT_bflike_LCDM_bf.o26551546.5 chi2=21113.587792 it=1721 t=5.421063s tot=2.591569h  errors=56
hlpTT_bflike_LCDM_bf.o26551546.6 chi2=20448.158418 it=3819 t=5.074077s tot=5.382750h  errors=57
hlpTT_bflike_LCDM_bf.o26551546.7 chi2=20448.667074 it=2867 t=6.624834s tot=5.275944h  errors=86
hlpTT_bflike_LCDM_bf.o26551546.8 chi2=20448.154056 it=3971 t=5.468623s tot=6.032194h  errors=86
hlpTT_bflike_LCDM_bf.o26551546.9 chi2=20448.225423 it=2925 t=3.550632s tot=2.884889h  errors=87
-------------------------------------------------------
chi2best= 20448.12723918  for hlpTT_bflike_LCDM_bf.o26551546.1
```

For each line you find : its name, $\chi^2_{\rm min}$, number of iterations, time per iteration, total time, number of (possible) `class` errors (warnings). So in our case, the best result was indeed for the first job, ie. the one analyzed in Sect. 3.2.1. Some "accidents" may happen (large χ^2 as for the 5th job) where some extreme random starting values did not allow a proper convergence. But basically we are confident the best fit initially found is indeed the best one.

3.2.3 The Hessian matrix

`MINUIT` allows to retrieve the Hessian matrix corresponding to the local curvature of the parameters around the minimum. It is produced when adding a name for the matrix to the arguments of the `Minimize` executable (which is the default in the `Minimise.sh` script) You should be careful when using it since:

- it is not very precise due to the numerical noise,
- it corresponds to second-order partial derivative matrix and cannot be identified to an "error matrix" in the case of strongly non linear parameters (James 1980).

It should therefore only be considered as indicative: the right way to obtain real "errors" is through the use of the profile-likelihood or MCMC methods. As we will see in Sect 3.4, it is however very convenient as the input proposal to the adaptive MCMC algorithm and we will check *a posteriori* that it was quite precise in this case (Fig. 11).

3.3 Profile-likelihoods

3.3.1 Properties

We now address the question of *interval estimation* for a single parameter (or a pair, rarely more) using the method of profile-likelihoods. Technically it consists in **fixing** values of a given parameter (let's call it θ_i) and performing a minimizations w.r.t all the others. For each θ_i scanned value, the $\chi^2(\theta_i) = \chi^2_{\rm min}$ value is reported on a graph that is finally offset to 0. We will focus our discussion on 1D profile-likelihoods, ie. when one scans a single parameter, but 2D ones (i.e scanning values over a grid) are also available in `CAMEL`.

The interest of reconstructing a profile-likelihood stems from the fact that by cutting it, one can reconstruct an α-level confidence intervals in the frequentist sense: if the experiment is repeated many times the interval covers the true values a fraction α of the times. We first consider the case of a Gaussian variable. Then $\Delta\chi^2$, considered as random variable of the data, is centered on the true value and follows a χ^2_1 distribution [4] (see eg. Severini (2000)). Then the

[4] or χ^2_2 for 2 parameters

dim	0.68	0.95	0.99
1 par	1	3.84	6.63
2 par	2.29	5.99	9.21

Table 1: Cutoff values to apply to $\Delta\chi^2$ for 1 or 2 parameters ("dim") in order to obtain 68,95 and 99% CL intervals

α-level interval is obtained by:

$$P(\Delta\chi^2) < \chi_\alpha \tag{1}$$

$$\text{where } \int_0^{\chi_\alpha} \chi_1^2(t)dt = \alpha. \tag{2}$$

We give in Table 1 some classical values of χ_α for 1 and 2D profile-likelihoods.

Although we built our discussion from the Gaussian case, it *actually holds for non-Gaussian cases too* (ie. non parabolic $\Delta\chi^2$ curves) and allows to reconstruct asymmetric intervals (w.r.t the MLE). As discussed in James (2007), in the general case one can always *imagine* a change of variable that would make the profile parabolic. The MLE invariance discussed Sect 3.2 more generally applies to likelihood ratios $\frac{\mathscr{L}(\theta_i)}{\mathscr{L}_{max}}$ or equivalently $\chi^2(\theta_i) - \chi^2_{\rm min}$. Thus the assertion made on the interval for the transformed variable ($\Delta\chi^2 = 1$) also applies to the original variable. Strictly speaking there can be a slight difference (because one transforms the data not the true value) but this only happens for very low statistics (a data samples) which is anyway always difficult to treat and a situation where the MLE is not necessarily optimal.

For completeness, we note that there are some cases where one wants to incorporate some physical bound (as for example a mass being always positive). This is most naturally performed in the Bayesian framework through the use of priors. The "unified approach" of Feldman & Cousins (1998) can be applied here and was demonstrated in Planck Collaboration Int. XVI (2014) and Henrot-Versillé et al. (2015).

Before showing a concrete example we would like to point out that because of the method, the best fit at the minimum of the $\Delta\chi^2(\theta_i)$ function is *the best-fit of the full set of parameter*s (the one obtained in Sect.3.2). This provides an interesting consistency test and ensures that, unlike in the MCMC approach, there is no difference when quoting the parameter value and the best fit solution.

3.3.2 Profile-likelihoods in practice

Building a profile-likelihood is then nothing more than running several minimizations fixing one parameter each time. Although there exist a `CAMEL` executable that loops on the variable and perform sequentially minimizations, one can take advantage of batch farms by sending each minimization on a separate worker. This can be performed with the `Profile.sh` script. To use it, one must first write a *scan-file* which specifies the profiled variable and its values. For instance, in a file `omega_b.scan` we put [5]

```
var="omega_b"
val=(0.022000 0.022020 0.022041 0.022061 0.022082 0.022102 0.022122 0.022143 0.022163 0.022184
0.022204 0.022224 0.022245 0.022265 0.022286 0.022306 0.022327 0.022347 0.022367 0.022388 0.022408
0.022429 0.022449 0.022469 0.022490 0.022510 0.022531 0.022551 0.022571 0.022592 0.022612 0.022633
0.022653 0.022673 0.022694 0.022714 0.022735 0.022755 0.022776 0.022796 0.022816 0.022837 0.022857
0.022878 0.022898 0.022918 0.022939 0.022959 0.022980 0.023000)
```

Then run:

```
Profile.sh path/to/hlpTT_bflike_LCDM.par /path/to/omega_b.scan
```

[5] this is only illustrative, you rarely need so many points. Note also that since this file is sourced by the script, one can use some `bash` commands (as loops) to generate these numbers.

When all the jobs finished, we concatenate the best-fit files [6] and represent the "chi2min" vs. "omega_b" variables reading the corresponding columns of the concatenated file (the profiled variable is always the column before the "chi2min" one). We obtained the points shown on Fig. 3.

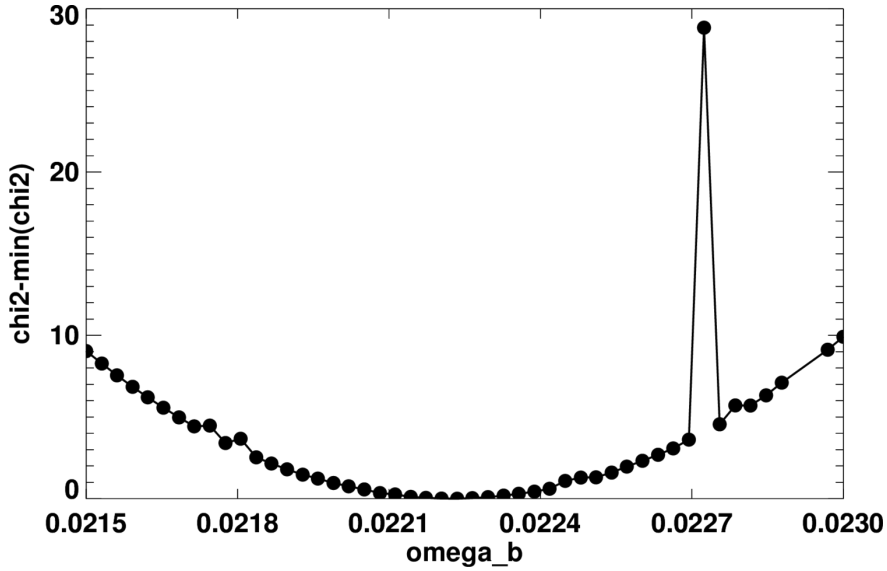

Figure 3: A first reconstruction of the omega_b profile likelihood points using `hlpTT_bflike_LCDM.par` with one single starting point

Obviously there has been a "glitch" near 0.0227 and some "wiggles" indicates the fit did not reached exactly its minimum in 3 or 4 places. As for the best fit case (Sect 3.2.1) to cure that issue we shoot different starting points. This does not need to be re-performed for each point, only the ones that are manifestly incorrect: this is the reason for writing explicitly the scan-file: you may re-edit it and keep only the values that needs to be improved.

Here we re-run (for the sake of illustration) on all the point (ie keep the same *scanfile*) shooting each time 6 different starting points for minimization:

```
Profile.sh path/to/hlpTT_bflike_LCDM.par /path/to/omega_b.scan 1-6
```

"1-6" means that jobs 1 to 6 will be launched for each scanned value. Each one (except for number 1) will shuffle randomly the starting values of each parameter in their specified [min,max] range. This shot is reproducible, so that if you need more starting points do not use "1-6" again but something as "7-10". This allows to increase gradually the statistics until reaching satisfactory continuity.

The script creates the `hlpTT_bflike_LCDM_mprof_omega_b/` directory and within it several other ones corresponding to the scanned values (which will contain each the outputs of the 6 jobs). Let us look at the outputs of all the jobs by concatenating all the `best_fit` files. All the points are shown on Fig. 4.

Some fits did not converge well, but what we are interested in is, for each scanned value, the minimal χ^2 value. There are in each case our "best best fit" solution and shown in red on Fig. 4. Using them the profile is now nicely continuous. For convenience, and since it is a common operation, we provide a script (in `work/tools/python`) to extract these points. Simply run it on the list of sub-directories containing the scanned values. Here we used (in the directory `hlpTT_bflike_LCDM_mprof_omega_b`)

```
python tools/python/bestbestfit.py omega_b*
```

where `omega_b*` is the pattern of the output directories. This creates the `profile_omega_b.txt` file.

[6] we will see later how to do that.

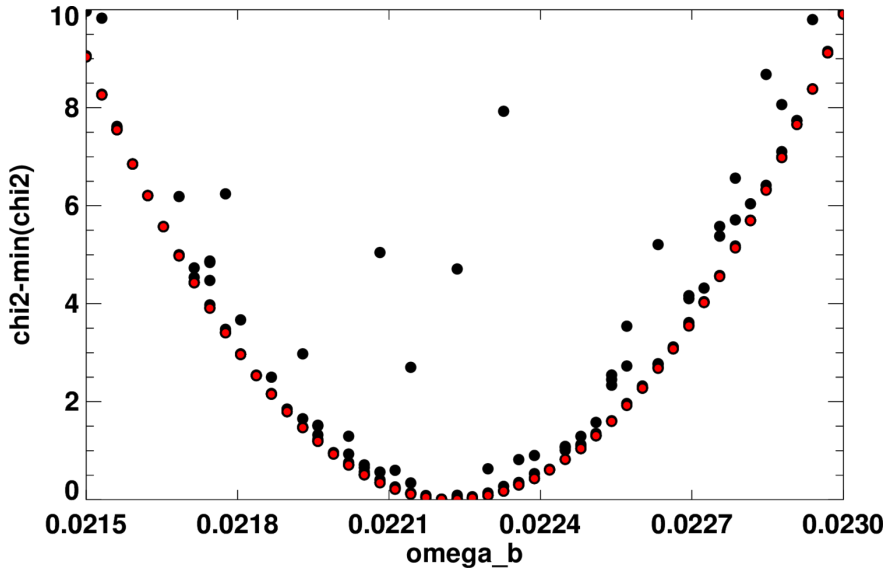

Figure 4: $\Delta\chi^2$minima reached with different starting points (6) for ω_b scanned values. The profile likelihood is obtained as the minimal values for each abcissa point and shown in red.

Then you interpolate smoothly between the profile points and determine the intersection with the $\Delta\chi^2 = 1$ line to obtain the 68% CL interval. You may quote "error bars" by comparing these limits to the position of the minimum. Be careful however to determine the minimum and limits on the interpolated curve not the data points. You can code this yourself in your preferred environment or use the `draw_profile.py` script which produced Fig. 5 running:

```
python tools/python/draw_profile.py profile_omega_b.txt
```

There are other ways to produce this plot (and others) through the `tools/python/camel.py` library (see the `CAMEL` web site for details).

There is actually more information we can extract from the set of best-fits. For each fixed ω_b value we have the multi-dimensional solution of the minimum. Then one can look not only at the $(\omega_b, \Delta\chi^2)$ curve but also at the best-fit solution achieved for the other variables ie. (ω_b, θ_i). This represents the covariance between ω_b and all the other variables (but not the *correlation* since we are missing σ_{θ_i}). These curves represents the (ω_b, θ_i) minimum valley of the likelihood and are similar (but not equivalent since there is no marginalization) to the major-axes of the 2D Gaussian posteriors in an sampling approach. In any case inspecting these "co-profile" is a good way to check that none of the parameters was blocked at some limit. We provide the python scritpt `coprofiles.py` to produce these plots. We ran:

```
python tools/python/coprofiles.py profile_omega_b.txt
```

and it produced the output shown on Fig. 6.

3.3.3 Profile-likelihood summary

The construction of a profile-likelihood may seem delicate because we entered in many details. In practice you should pay attention to the following points:

- always use a precision file (with `class`). the default `hpjul2.pre` should be sufficient in most cases.
- choose your scanning points carefully. You need to focus only on the $\Delta\chi^2 \lesssim 5$ region (which may require some iterations): it is useless (and difficult) to regularize points outside.

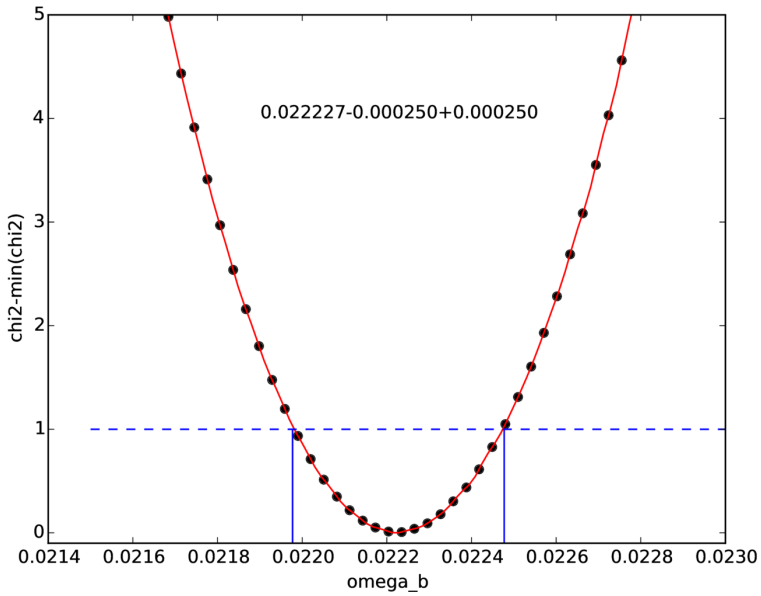

Figure 5: Construction of the 68%CL interval by thresholding the interpolated ω_b profile-likelihood (in red) at 1.

- although we described the method with many points (50), you don't need as much to construct a reliable curve (10 to 20 are sufficient).
- you should first try with the parameter-file fixed starting point (ie. range 1-1 in `Profile.sh` which is the default)
- then if you notice some continuity problem
  - if there are little, simply discards those points before performing a smooth interpolation among the remaining "good" ones.
  - if there are many, try more starting points on the problematic values
- once you are happy with your profile, always check the "co-profile" to see if some parameters are not blocked by the [min,max] bounds.

3.3.4 Results on ΛCDM cosmological parameter

We apply this methodology to the remaining 5 cosmological parameters in the file `hlpTT_bflike_LCDM.par` and obtain the profile-likelihoods shown on Fig. 7

3.4 MCMC

3.4.1 Real-life Metropolis

The Monte-Carlo Markov Chain method draws samples according to any distribution (for a nice introduction see eg. Bardenet (2013)). When applied to a likelihood function, it may serve to reconstruct the unknown parameters distribution (posteriors), a concept not accepted in a classical "frequentist" perspective. According to Bayes theorem:

$$p(\theta|x) = \frac{p(x|\theta)\pi(x)}{p(x)} \propto \mathscr{L}(\theta|x)\pi(\theta) \tag{3}$$

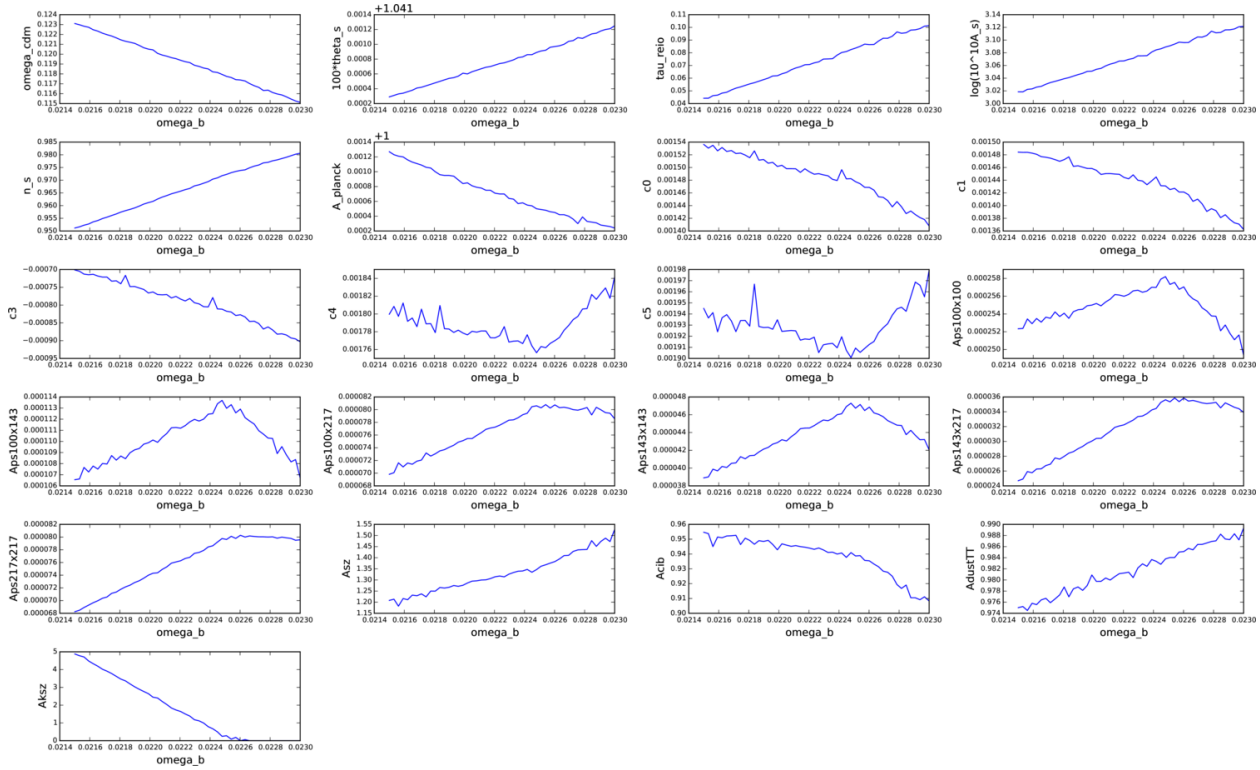

Figure 6: Each plot shows the value of the θ_i parameter w.r.t to the scanned variable (ω_b) at the best-fit solutions. We call these curves "co-profiles". They are an indication whether the 2D profile-likelihoods are Gaussian (in which case they should be slopes).

$\mathscr{L}$ being the likelihood of the data and π the prior distribution that reflects the degree of belief on the parameter before considering the experiment. The latter is rarely discussed (ie. considered as equal to 1 in some reasonable range although this is not necessarily the most non-informative prior (Jeffreys 1983)) in which case the likelihood of the parameters equals their probability. Sampling from the likelihood then gives directly the parameters posterior distribution.

The founding father of all MCMC algorithms is the Metropolis one (Metropolis & Ulam 1949; Metropolis et al. 1953). Its mathematical properties have been studied for decades and it is certainly the best characterized one (Roberts & Casella 1998; Gilks et al. 1996). Although it is *mathematically* guaranteed to converge to the proper distribution at infinity, this is often too far in real case studies, in particular in cosmology, where Boltzmann solver computations are slow. In practice it means that the *proposal distribution* used to perform the random jumps, generally a central multivariate Gaussian with some covariance matrix, must be known precisely. This matrix is often complicated because of important off-diagonal elements that can hardly be a priori inferred (see eg. Fig. 11): not incorporating them will not allow a proper sampling of the likelihood on a reasonable timescale. Therefore the matrix is reconstructed iteratively from the empirical covariance of the samples but it is a very time-consuming process. *Reproducing* results when the proposal is *given* is simple but hides the lengthy previous steps. A soon as one want to study new variables/likelihoods the proposal matrix must be rebuild again.

We have therefore implemented in `CAMEL` an adaptive scheme of the Metropolis algorithm. In this approach the proposal covariance matrix is built on the fly: the adaptation level is gradually decreased so that one ends up with a fixed matrix and then follows the classical Metropolis algorithm. The initial matrix does not need to be known precisely and can be obtained in 2 ways:

1. one may use the Hessian matrix obtained from minimization (Sect 3.2.3).

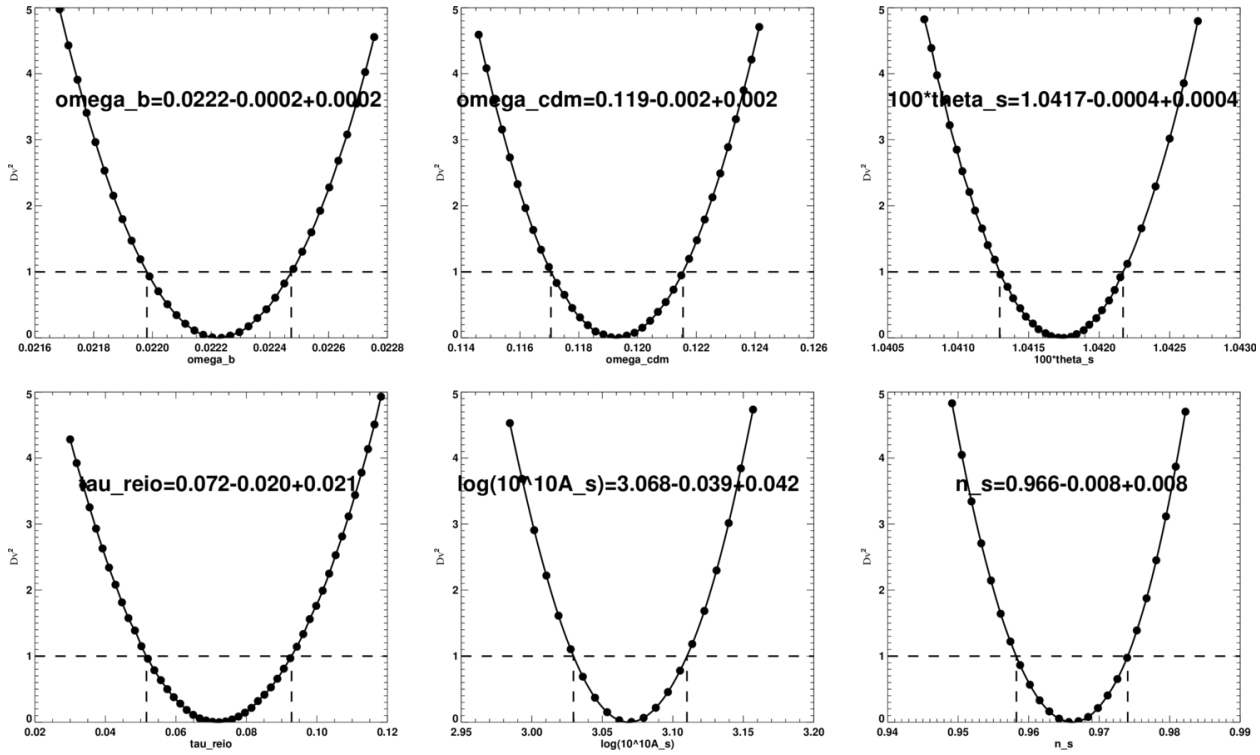

Figure 7: Profile-likelihood reconstructed for the 6 ΛCDM cosmological parameters with the `hlpTT_bflike_LCDM.par` parameter file. The line at 1 gives the 68% CL region of each parameter.

2. often one just want to expand a set of parameters, and already knows from the empirical covariance a precise matrix of a subset of parameters: then just add a crude estimate of the new parameters errors on the diagonal and let the algorithm do the work.

3.4.2 Adaptive Metropolis (AM)

The original algorithm (Haario et al. 2001) works in the following way:

1. Choose a starting point X_0, a starting covariance matrix Σ_0 and a tuning parameter c;
2. generate a value Y from a proposal density $N(X_{t-1}, c\Sigma)$;
3. evaluate the test ratio $\alpha = min[1, \frac{\pi(Y)}{\pi(X)}]$;
4. generate a value u uniformly distributed in $[0, 1]$;
5. if $u \leq \alpha(X_t, Y)$ set $X_{t+1} = Y$ else set $X_{t+1} = X_t$;
6. update running mean and covariance:

$$\mu_t = \mu_{t-1} + \frac{1}{t}(X_t - \mu_{t-1}), \quad \Sigma_t = \Sigma_{t-1} + \frac{1}{t}((X_t - \mu_t)(X_t - \mu_t)^T - \Sigma_{t-1}) \tag{4}$$

7. Increment t.

In practice, this algorithm requires developments in order to be used in non-trivial conditions. If we try to start the adaptation from the very first steps, the correction we make to the initial guess matrix Σ_0 is too big and numerically unstable. We thus introduced the parameter t_0, that can be tuned, and that represents the number of steps we wait before starting the adaptation. We use as default $t_0 = 2000$ steps which gave us good results in all the cases we studied (from N=20 to 50 parameters). If you are pretty sure of your initial matrix you may lower this number. Anyway the algorithm is very weakly dependent on this choice. In the $t \leq t_0$ phase the chain starts exploring the parameter space, and it is not essential to have a very reliable first estimate for the covariance matrix; it is preferable to have a high acceptance rate and underestimate the optimal step than to risk remaining stuck at the starting point. The scale c in the proposal is thus chosen to be quite small (default 0.001)

Then at step t_0 we calculate the sample variance of the chain. The estimate of the covariance between the parameter j and the parameter k is simply

$$\Sigma_{jk} = \frac{1}{t_0 - 1} \sum_{i=1}^{t_0} (X_{ij} - \bar{X}_j)(X_{ik} - \bar{X}_k). \tag{5}$$

This reconstructed matrix is a fair first guess of the covariance of the posterior. Hence, the scale c is now set to the optimal value (for Gaussians) of $2.4^2/d$ where d is the number of parameters (Gelman et al. 1996; Dunkley et al. 2005) and we start running the matrix adaptation.

Since the likelihood is not necessarily Gaussian and the matrix not perfect (otherwise we won't need to run!), we introduced another refinement that concerns the scale parameter c. At a time t_s we begin to adapt the scaling parameter according to

$$c_{t+1} = c_t \cdot (1 - \frac{1}{t}) + \frac{1}{t}(\text{a.r.} - 0.25) \tag{6}$$

where "a.r" is the acceptance rate (ie. the number of accepted moves over the total) and is computed with the previous 100 samples. The purpose of this stage is to tune the scale parameter c in order to reach an a.r of about 0.25, a classical "rule of thumb" for the Metropolis algorithm, and that gave indeed in our experience the best results. The t_s default in `CAMEL` is 10000 samples. You may change this number (with counts steps after t_0) but note that the $\frac{1}{t}$ factors in Eq. 6 makes the scale factor converge rapidly to a constant and we end up with a classical Metropolis-Hastings algorithm.

Keeping memory of all its history, the process is no more Markovian. So, does the adaptive algorithm have the correct ergodicity properties? It is indeed mathematically difficult to deal with non basic MCMC methods, and a lot of the algorithm used by the scientific community are not assured to have the right properties for convergence. However in the Adaptative-Metropolis case, the convergence was proven (Andrieu & Moulines 2006) because the asymptotic dependency between the elements of the chain is weak enough to apply large number theorems.

3.4.3 Runing AM

Let us see how to use the AM. One needs to specify a few extra information to the parameter file. Most have correct default values but you should at least define the number of samples ("length") and the path to the initial proposal covariance matrix ("proposal_cov"):

```
length=100000  # specify the chain length
bunchSize=1000 # the chain is dumped every bunchSize
algo=ada       # default algorithm (or "metro" for standard Metropolis)
proposal_cov=/path/to/your/covmatrix.dat # path to your first guess covariance matrix
t0=2000        # number of steps for initial cov construction
ts=10000       # number of samples over t0 before starting the scale adapation
scale=0.001    # initials scaling (shoud be low)
#seed=12345     # seed for the random number generator for proposal: if not defined drawn
    randomly
#seedU=78776    # seed for uniform shots in algorithm: if not defined drawn randomly
```

Some important remarks

- the program will dump the data into a specified text file every `bunchSize` iterations. This is interesting for running remotely (as in batches) and flushing the result through the network without too intensive IOs. You can then investigate results without waiting the end of the run.

- there are 2 seeds for the random number generators. **If you fix these numbers you will recover the same results** which is probably not what you want if you run multiple chains (but is interesting for reproducibility). You can then implement some logic to draw some random numbers for these seeds, or do not specify them at all: then they will be drawn randomly from the machine environment.
- the covariance matrix should be given as a text file. You can (should?) use the Hessian matrix of the minimization (see Sect.3.2.3). If not specified a diagonal matrix will be built from the parameter file "errors" but it is probably a bad idea to do so...
- because of the hopping nature of the algorithm, it is unnecessary (and CPU consuming) to use a high precision file. Don't put any `precisionFile` line in your parameter file: this will use `class` default settings which are sufficient.

If you are using a full Boltzmann solver (as `class`) you may also want to add to the output some *derived* parameters which are computed at each sampled step. The list of them can found in the `src/camel/Engine.hh::get` method. You add to your parameter file some lines as:

```
derived=H0
derived=z_reio
derived=YHe
derived=sigma8
do_mPK=true
```

(note that you need to activate the "do_mPk" flag for sigma8 (which turns on the matter power spectrum computations)

You may also add some redshift dependent values as

```
derived=sigma8(0.57)
derived=f(0.57)
class z_pk 0,0.57
```

The last line is here to ask the exact computation of P(k) at z=0 and 0.57.

In order to use a script to submit jobs with AM, just put the parameters discussed in this section in some text file (for instance `mcsetup.txt` and run:

```
runMCMC.sh /path/to/parfile mcsetup.txt
```

This will launch 4 openMP jobs (you can change this number by redefining the NCHAINS environment variable). The chains will be written as plain ASCII files under the name `samples*.txt` in the directory created.

3.4.4 Convergence

We discuss here the question of MCMC convergence. We illustrate it with the chains produced with the `hlpTT_bflike_LCDM.par` run and with the settings discussed in the previous part. Note that we used directly for the initial covariance matrix the Hessian obtained in Sect 3.2.1. We ran 4 chains of 150000 samples each which produced 4 output files.

There exists no exact mathematical way to determine when (if) the chain reached "convergence", ie. entered a regime where it samples correctly the likelihood (in all dimensions). But there are some tricks and we discuss here two classical ones.

The first thing you should look at is simply each variable evolution ("trace plots"). This is shown below for one chain (obviously all should be looked at).

In this exercise for example, the "Aszxcib" variable began to be widely explored only lately (after $\simeq$50000 samples). This impacted the "Acib and all the "Aps" variables evolution.

This kind of plot only tests the *stationarity* of the chains. The real challenge however is to reach *ergodicity* (which requires stationarity) meaning that the samples can be used to compute statistical expectation values (ie. "time" integrals can be used to replace "ensemble" averages) This is the idea behind the Gelman-Rubin (GR)test Gelman & Rubin (1992) which is run on a few chains and compares the inter- and cross- variance of the samples. It produces R_i ,the "potential scale reduction factor" for each variable and the usual prescription is to assume that the variables reached ergodicity when all $R_i - 1 \leq 0.03$ (but you may feel more comfortable with 0.01).

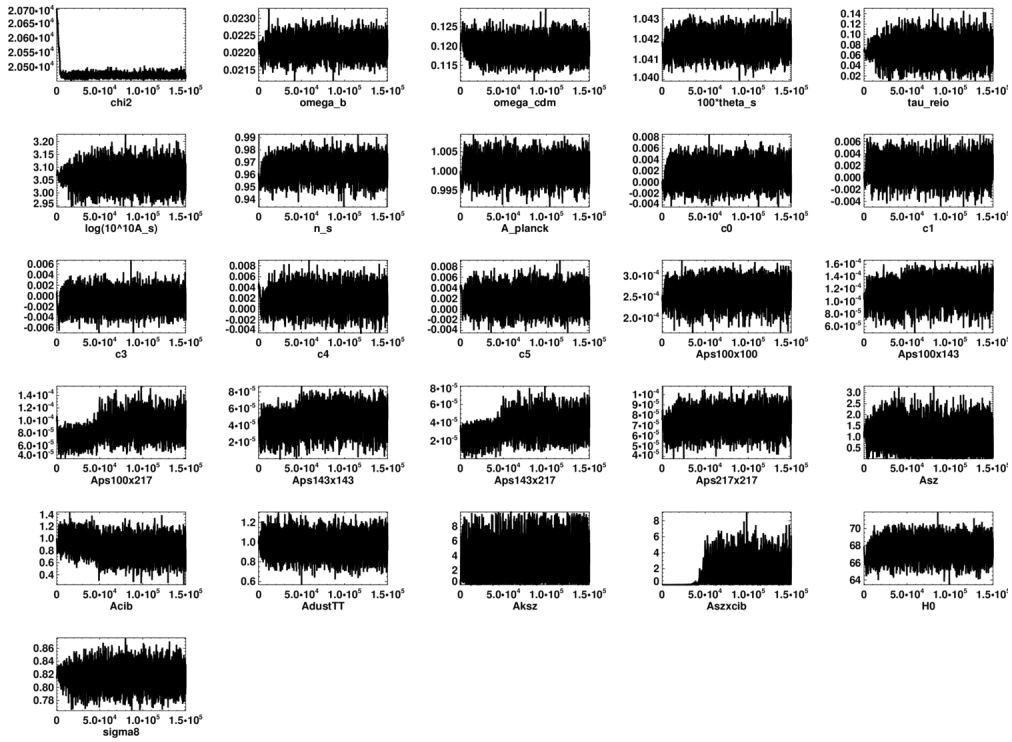

Figure 8: Plots of the values of each variable generated with the MCMC algorithm for one chain using `hlpTT_bflike_LCDM.par`. The first one displays the χ^2 evolution. The last 2 variables are derived parameters.

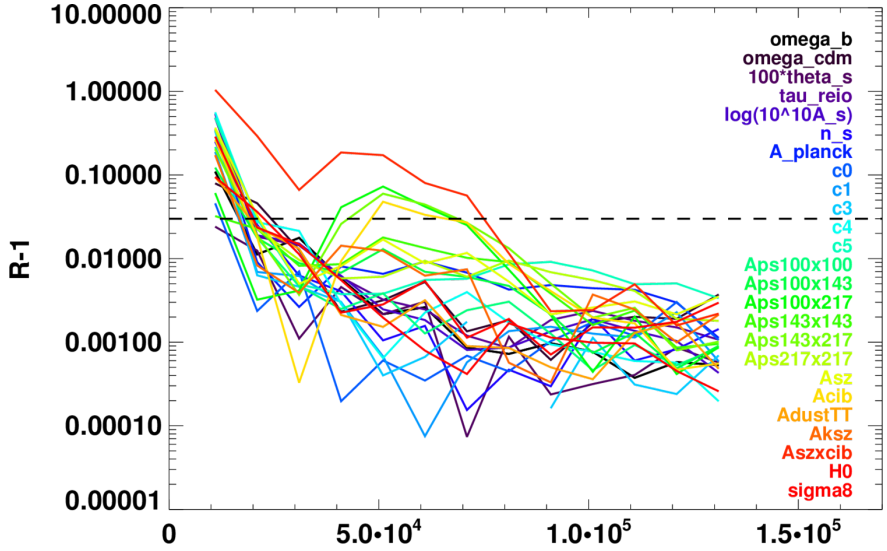

Figure 9: Gelman-Rubin test on the 4 chains for the sampled `hlpTT_bflike_LCDM.par` file. The standard cut is to use samples above $(R-1) \geq 0.03$ and is shown as a dashed line.

The result in our example, as implemented in `CAMEL` in `python` or `idl`, is shown on Fig. 9.

The GR-test confirms that "Aszxcib" began to be sampled correctly only after $\gtrsim$ 80000 steps. Once its sampling begun ($\gtrsim$ 50000) it indeed affected the "Aps" and "Acib" variables which is reflected in some increase of their R values before finally converging to low values.

While a number of other tests exists (we also implemented the Geweke one), GR produces generally satisfactory results (for simple posteriors).

3.4.5 MCMC analysis

Given the GR-test, we decided to keep samples above 100000 steps and concatenate them into a single "well-converged" chain.

1D posterior distributions are then obtained by histograming (and possibly smoothing) each variable individually . They are shown on Fig. 10. For further discussion, we also superimpose the best-fit values with Gaussian errors taken from the diagonal of the Hessian matrix.

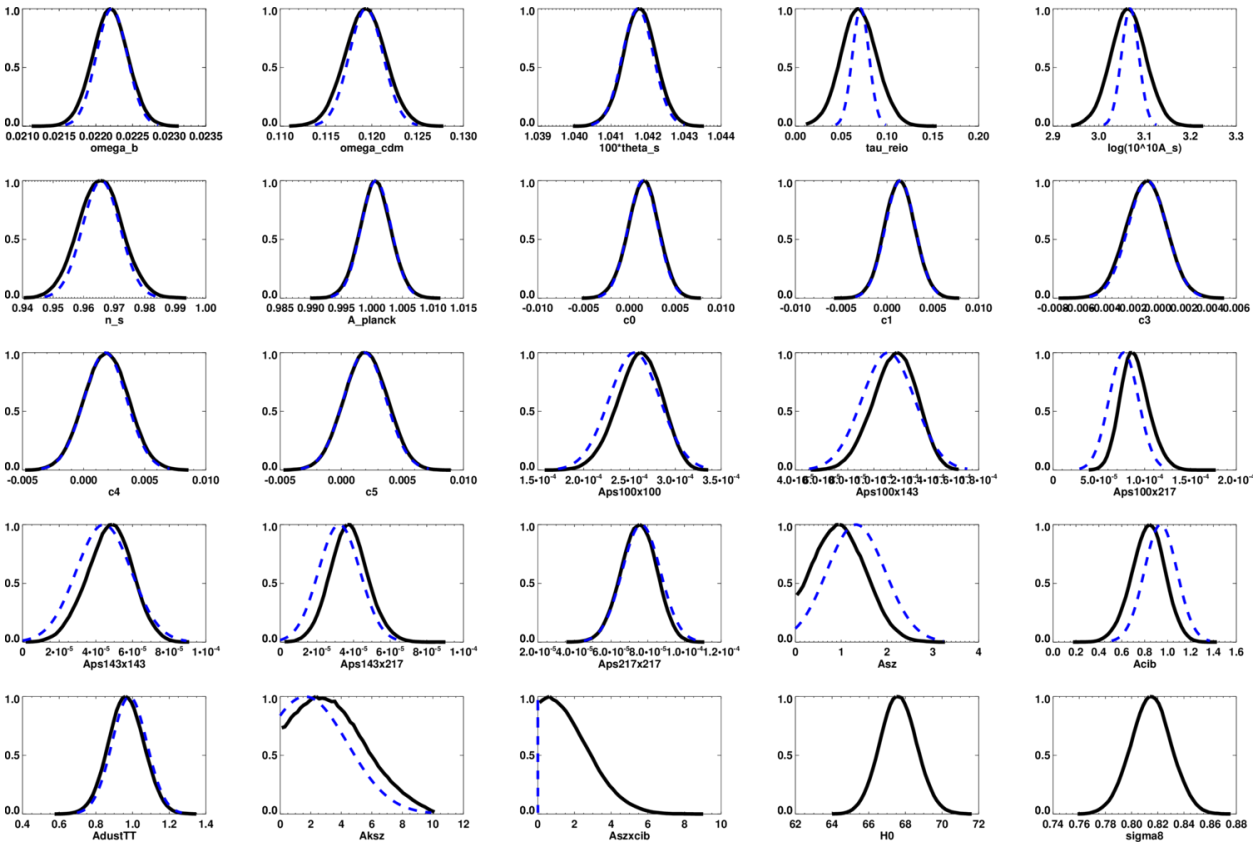

Figure 10: full lines: 1D posterior distributions of MCMC samples obtained by histograming + smoothing the well-converged chain. For clarity the maximum is normalized to 1. We also superimpose in dashed the output of the best-fit values with errors taken from the Hessian matrix. The last 2 parameters are MCMC-derived ones and have thus no associated best-fits.

First we note that the best-fit positions (dashed curve peaks) sometimes do not fall on the maximum of the posteriors (the "mode"). This is known as "volume effects" and will be discussed in Sect. 4.2. Then the width of the distributions as inferred from the Hessian diagonal elements are quite good, but for "AszxCIB" which was stuck at 0

during the minimization. This explains why it took some time for this parameter to be well sampled as was singled out by the GR-test. Despite this very wrong input, the algorithm still converged in a reasonable time.

Let us now study the correlations among the parameters. One can compute the empirical covariance matrix between the estimated parameters

$$\hat{C} = \frac{1}{N-1} \sum_{i=1}^{N} \Delta_i^T \Delta_i \tag{7}$$

where the i^{th} realization deviation vector is $\Delta_i \equiv \vec{p}_i - \hat{\mu}$, $\vec{p}$ being the vector of parameters and $\hat{\mu}$ its statistical mean. We convert it to a correlation matrix and show it on the left of Fig. 11 (we drop the "H0" and "sigma8" parameters which are derived parameters).

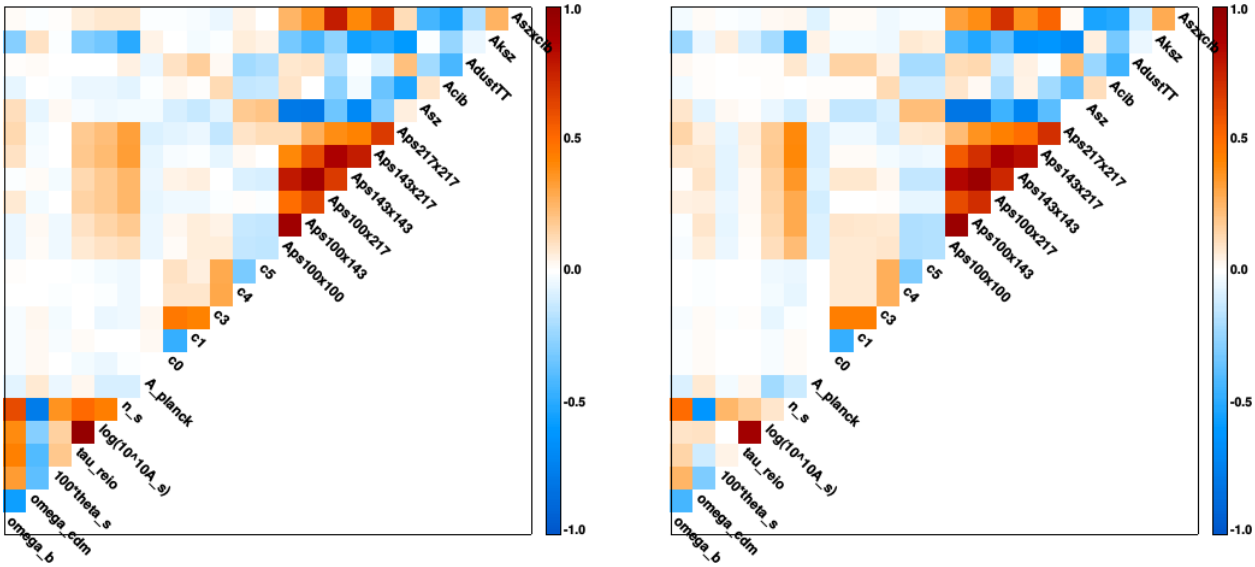

Figure 11: Correlation matrix between all the parameters estimated from the sampled posterior (on the left) and the Hessian matrix (on the right).

The comparison with the Hessian matrix on Fig. 11(b) reveals a posteriori that `MINUIT` performed quite well in determining the correlation structure. However note that the bottom part (cosmological parameters) is not perfect.

A famous representation in a single figure of both the 1D and 2D posterior distributions is in the form of a "triangle plot" available in `CAMEL` through the `camel.py` library or through the `GetDist` package (that can read `CAMEL`'s chains) and is shown on Fig. 12.

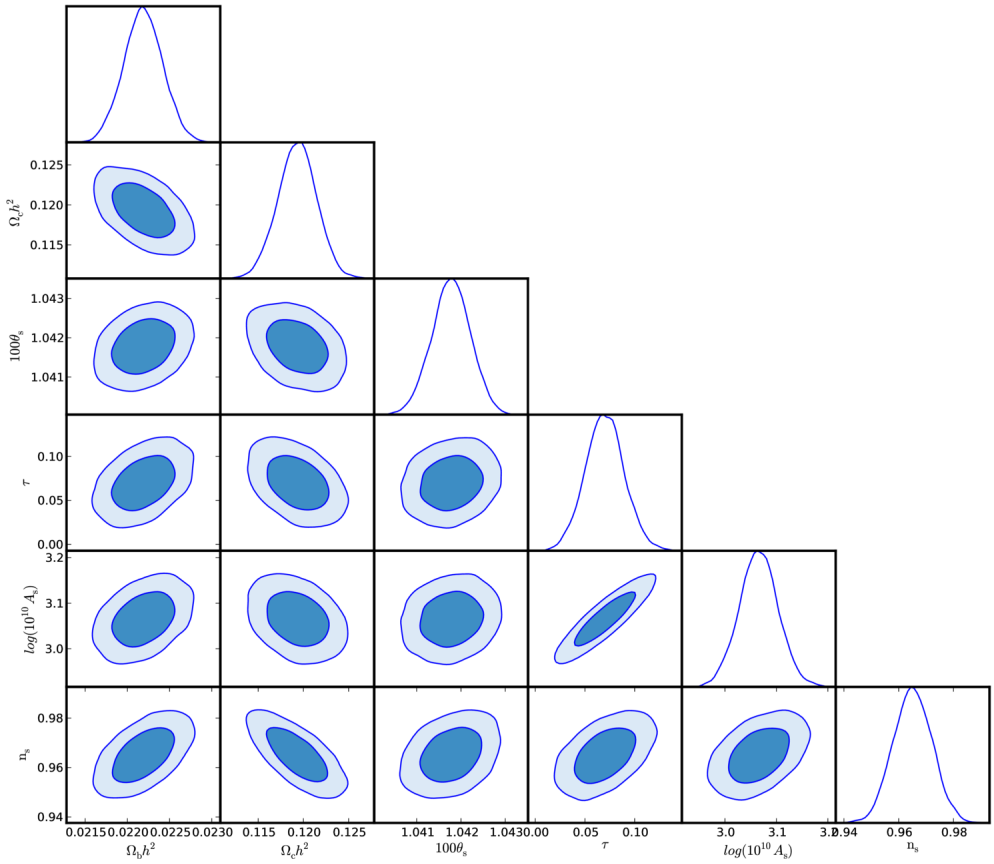

Figure 12: "Triangle plot" for the MCMC samples showing the 1D and 2D posterior distributions of the cosmological parameters from the `hlpTT_bflike_LCDM.par` run. Ellipses show the 68 and 95% credible intervals.

4 Comparison of the results

4.1 ΛCDM parameters

We can now check if *in this working case* both statistical methodologies lead to the same intervals. We show on Fig 13 the comparison between the 1D posteriors and the profile-likelihoods of the cosmological parameters as obtained in the previous sections. For this latter and the sake comparison, we represent $e^{-\Delta\chi^2/2}$ and the $\Delta\chi^2 = 1$ cut translates into a $e^{-1/2} \simeq 0.6$ one.

We therefore see that for ΛCDM using the *Planck* `Hillipop+bflike` likelihoods both the Bayesian and frequentist approaches give very similar results on cosmological parameters, a result coarsely shown with other *Planck* likelihoods (Planck Collaboration Int. XVI 2014; Planck Collaboration XIII 2015).

One should not conclude however it is a general statement: it happens here because the data constrain well all the parameters, in particular the use of the low-ℓ likelihood breaks the $(\tau, \ln(10^{10}A_s))$ degeneracy from the high-ℓ part. The posteriors are close to Gaussian so that their 1σ interval corresponds neatly to the one obtained with the profile-likelihoods cut at $e^{-1/2}$. More generally, for a single parameter, the Bayesian credible interval will be a genuine "confidence level" if and only if the parameter is a "location" one, ie. $\mathscr{L}(x;\theta) = \mathscr{L}(x-\theta)$ or some transformation of the likelihood leads to it (Porter 1996). In the multi-dimensional case, some differences may appear due to the Bayesian marginalization. This is called the "volume effect" and is discussed next.

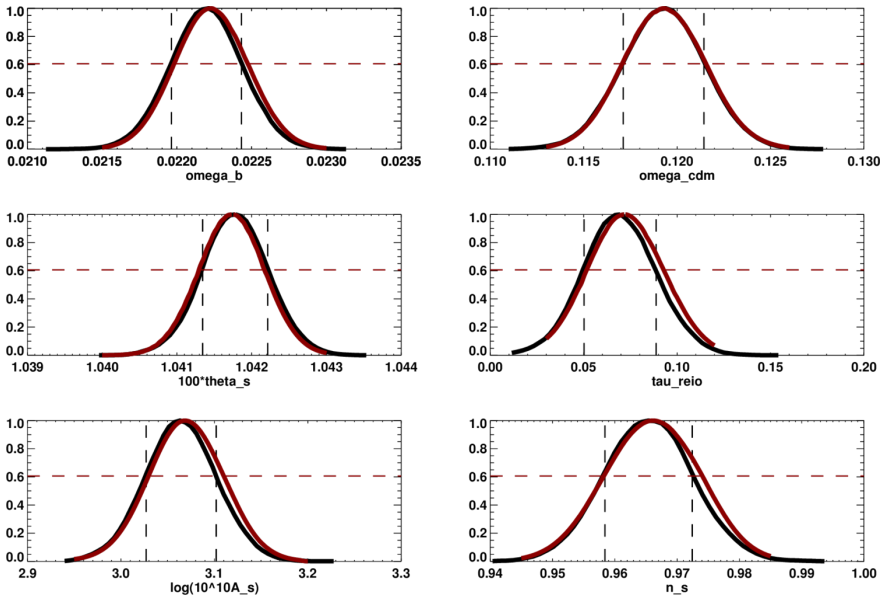

Figure 13: The MCMC posterior distributions (in black) as shown previously on Fig. 10 are compared to the profile-likelihoods (in red , same than those on Fig. 7 but represented here as $e^{-\Delta\chi^2/2}$) for the ΛCDM cosmological parameters using in both cases the same `hlpTT_bflike_LCDM.par` setup. The dashed lines are here to recall that 68% level interval are not obtained in the same way: for posteriors the interval is given on the abscissa (as mean $\pm 1\sigma$), while for profile likelihood the prescription $\Delta\chi^2 = 1 \implies \mathscr{L}/\mathscr{L}_{\rm max} = e^{-1/2}$ defines the cut along the ordinate. Both intervals agree in the Gaussian case.

4.2 Likelihood volume effects

In order to illustrate the difference between the profile-likelihoods and posterior distributions with several parameters, we begin with a toy example. Suppose the likelihood has 2 parameters (x_1, x_2) and that its shape is of the type shown on the upper plot of Fig. 14. Then the marginalized posterior for x_1 reads

$$p(x_1) \propto \int \mathscr{L}(x_1, x_2) dx_2, \tag{8}$$

while the profile-likelihood is

$$\mathscr{L}(x_1) \propto \max_{x_2} \mathscr{L}(x_1, x_2). \tag{9}$$

One can then easily understand why the maximum of the profile-likelihood lies at the same position than the 2D maximum. This is not necessarily the case for the marginalized posterior where some part of the "volume" of the 2D likelihood may shadow the projection. Indeed these methods answer different questions: in the case of the posterior "what is the value of x_1 given the possible x_2 values?", while for profile-likelihood, "what is the value x_1 for the most likely x_2 values?".

We now give an example where this situation happens. It was shown in Couchot et al. (2015) that when removing the low-ℓ part of the *Planck* baseline likelihoods and keeping only the high-ℓ one (named `Plik`), the profile-likelihood method leads to an optical reionization depth of:

$$\hat{\tau} = 0.17 \pm 0.04 \ (68\%, \text{Planck-high-}\ell\text{/profile-likelihood}). \tag{10}$$

The τ posterior distribution, obtained by running the AM method, is shifted to lower values is as shown on Fig.15. The peak of the distribution (the "mode") is located around 0.15 , [7]. By integrating the tails, one can reconstruct the

[7] and even slightly less for the mean and median

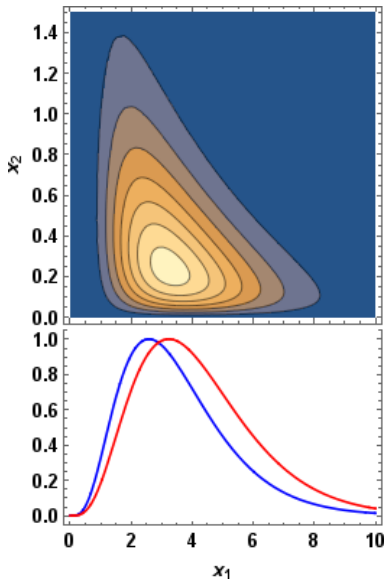

Figure 14: A toy example of a 2D likelihood (upper plot) and (bottom plots)the corresponding distributions for the x_1 marginalized posterior (in blue) and its profile-likelihood (in red).

68% central credible interval, which gives $[0.11, 0.18]$ and taking the mode as the best estimate, we obtain:

$$\hat{\tau} = 0.15^{+0.03}_{-0.04}\ (68\%, \text{Planck-high-}\ell\text{/posterior}). \tag{11}$$

The choice between both methodologies can play a role. For instance, from an LFI-based low-ℓ analysis (Planck Collaboration XI 2015), *Planck* reports a value of [8]: $\hat{\tau} = 0.067^{+0.023}_{-0.021}(68\%, \text{Planck-LFI low-}\ell)$. It is in 2.2σ tension with Eq.10 but 1.8σ with Eq.11. Not over-emphasizing the 2σ boundary, the use of the Bayesian result may have hidden this tension that is now increased by the new HFI-based result (Planck Collaboration Int. XLVII 2016): $\hat{\tau} = 0.058 \pm 0.012, (68\%, \text{Planck-HFI low-}\ell)$ to respectively 2.6σ (profile) and 2.2σ (posterior).

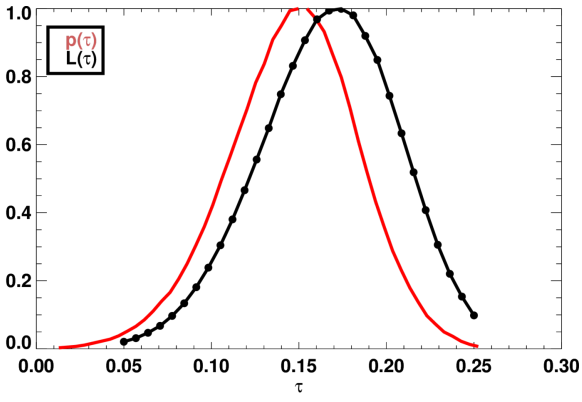

Figure 15: Reionization optical depth posterior (in red) obtained by sampling, with the AM algorithm, the *Planck* high-ℓ-only likelihood (`plik`). The corresponding profile likelihood is superimposed in black.

[8] This result is independent of the methodology used

5 Summary

We have seen how comparing different statistical approaches can be fruitful in cosmological analyzes. A possible workflow can be:

1. first run an accurate Minimization with `MINUIT` with a few starting points. Obtain the best-fit and the Hessian matrix.
2. use the Hessian as the starting matrix to the proposal of the Adaptive Metroplis MCMC algorithm.
3. if some posterior modes differ significantly from the best-fit value, run single profile-likelihoods on those parameters.
4. if some parameters are of particular importance (as (w_0, w_a) in future surveys [9]), always run a profile-likelihood on them to verify the influence of priors and volume effects in the MCMC run.

`CAMEL` provides the tools to perform such rich studies *using the very same input file*. See the `camel.inp3.fr` web site for more information. The software is collaborative and contributions are most welcome.

References

Allison, R. & Dunkley, J. 2014, MNRAS, 437, 3918, arXiv:1308.2675

Anderson, L., Aubourg, E., Bailey, S., et al. 2014, Monthly Notices of the Royal Astronomical Society, 441, 24

Anderson, L., Aubourg, E., Bailey, S., et al. 2012, Monthly Notices of the Royal Astronomical Society, 427, 3435

Andrieu, C. & Moulines, E. 2006, Ann. Appl. Probab., 16, 1462

Audren, B., Lesgourgues, J., Karim, B., & Prunet, S. 2013, JCAP, 1302, 001, arXiv:1210.7183

Bardenet, R. 2013, EPJ Web of Conferences, 55, 02002

Betoule, M., Kessler, R., Guy, J., et al. 2014, Astronomy & Astrophysics, 568, A22

Beutler, F., Blake, C., Colless, M., et al. 2011, Monthly Notices of the Royal Astronomical Society, 416, 3017

Blake, C., Davis, T., Poole, G. B., et al. 2011, Monthly Notices of the Royal Astronomical Society, 415, 2892

Blas, D., Lesgourgues, J., & Tram, T. 2011, J. Cosmol. Astropart. Phys., 2011, 034

Christensen, N., Meyer, R., Knox, L., & Luey, B. 2001, Classical and Quantum Gravity, 18, 2677, arXiv:astro-ph/0103134

Couchot, F., Henrot-Versillé, S., Perdereau, O., et al. 2015, ArXiv e-prints, arXiv:1510.07600

Das, S., Louis, T., Nolta, M. R., et al. 2014, J. Cosmol. Astropart. Phys., 2014, 014

Davidon, W. & Laboratory, A. N. 1959, Variable metric method for minimization, AEC research and development report (Argonne National Laboratory)

Dunkley, J., Bucher, M., Ferreira, P. G., Moodley, K., & Skordis, C. 2005, Monthly Notices of the Royal Astronomical Society, 356, 925

Feldman, G. J. & Cousins, R. D. 1998, Phys. Rev. D, 57, 3873, arXiv:physics/9711021

Fendt, W. A. & Wandelt, B. D. 2007, ApJ, 654, 2

[9] 2D profile-likelihoods are also implemented

Feroz, F., Hobson, M. P., & Bridges, M. 2009, MNRAS, 398, 1601, arXiv:0809.3437

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306, arXiv:1202.3665

Gelman, A., Roberts, G. O., & Gilks, W. R. 1996, In J. M. Bernado et al., Oxford University Press, 599

Gelman, A. & Rubin, D. B. 1992, Statistical Science, 7, 457

George, E. M., Reichardt, C. L., Aird, K. A., et al. 2015, ApJ, 799, 177

Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. 1996, Markov Chian Monte Carlo in practice (Chapmanand Hall/CRC)

Goodman, J. & Weare, J. 2010, Comm. App. Math. and Comp. Sci, 5

Haario, H., Saksman, E., & Tamminen, J. 2001, Bernoulli, 7, 223

Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015, Monthly Notices of the Royal Astronomical Society: Letters, 450, L61

Henrot-Versillé, S., Robinet, F., Leroy, N., et al. 2015, Class. Quantum Grav., 32, 045003

Hobson, M. P., Jaffe, A., Liddle, R., Mukherjee, P., & Parkinson, D. 2010, Bayesian methods in cosmology (Cambridge University Press)

James, F. 1980, Computer Physics Communications, 20, 29

James, F. 2007, Statistical Methods in Experimental Physics (World Scientific)

James, F. & Roos, M. 1975, Comput. Phys. Commun., 10, 343

Jeffreys, H. 1983, Theory of Probability (3d ed) (Oxford University Press)

Lewis, A. & Bridle, S. 2002, Phys. Rev. D, 66, 103511, arXiv:astro-ph/0205436

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, The Journal of Chemical Physics, 21, 1087

Metropolis, N. & Ulam, S. 1949, Journal of the American Statistical Association, 44, 335

Padmanabhan, N., Xu, X., Eisenstein, D. J., et al. 2012, Monthly Notices of the Royal Astronomical Society, 427, 2132

Planck Collaboration Int. XLVII. 2016, A&A, submitted, arXiv:1605.03507

Planck Collaboration Int. XVI. 2014, A&A, 566, A54, arXiv:1311.1657

Planck Collaboration XI. 2015, ArXiv e-prints, arXiv:1507.02704

Planck Collaboration XIII. 2015, ArXiv e-prints, arXiv:1502.01589

Planck Collaboration XV. 2016, A&A, in press, arXiv:1502.01591

Porter, F. 1996, Nuclear Instruments and Methods in Physics Research A, 368, 793

Reichardt, C. L., Shaw, L., Zahn, O., et al. 2012, ApJ, 755, 70

Roberts, C. P. & Casella, G. 1998, Monte Carlo Statistical Methods (Springer)

Severini, T. 2000, Likelihood Methods in Statistics, Oxford science publications (Oxford University Press)

Skilling, J. 2004, AIP Conference Proceedings, 735, 395

Story, K. T., Reichardt, C. L., Hou, Z., et al. 2013, ApJ, 779, 86

Tristram, M., Macias-Perez, J. F., Renault, C., & Santos, D. 2005, Monthly Notices of the Royal Astronomical Society, 358, 833

Zuntz, J., Paterno, M., Jennings, E., et al. 2015, Astronomy and Computing, 12, 45, arXiv:1409.3409

Appendix A : a (very) complete parameter file

The following file shows the complete list of likelihoods with their associated nuisance parameters. The Engine used is `class`.

```
#####################################################################
#Boltzman Engine
engine=class

#class setup
class      N_ncdm                 1
class      m_ncdm                 0.06
class      N_eff                  2.046

class      k_pivot                0.05
class      lensing                yes
class      sBBN\ file             bbn/sBBN.dat

precisionFile=class_pre/hpjul2.pre
#####################################################################



#####################################################################
#Cosmological model
par    omega_b              cosm     0.02224      0.00027    0.017    0.027
par    omega_cdm            cosm     0.1192       0.0026     0.09     0.15
par    100*theta_s          cosm     1.0418       0.6E-04    1.03     1.05
par    tau_reio             cosm     0.07         0.13E-01   0.01     0.20
par    log(10^10A_s)        cosm     3.07         0.025      2.7      3.5
par    n_s                  cosm     0.96         0.0070     0.9      1.1
#####################################################################



#####################################################################
#LIKELIHOODS

#-----------------------------
#CMB low-l (BFLIKE)
clikfile=planck_data/low_l/bflike/lowl_SMW_70_dx11d_2014_10_03_v5c_Ap.clik
#-----------------------------


#-----------------------------
#CMB high-l (HILLIPOP)
HiLLiPOP=HiLLiPOP/DX11dHM_superExt_CO_TT.lik
par    A_planck             nui      1            0.001        0.9          1.1
par    c0                   nui      0.           0.001      -0.05    0.05
par    c1                   nui      0.           0.001      -0.05    0.05
fix    c2                   nui      0.           0.001      -0.05    0.05
par    c3                   nui      0.           0.001      -0.05    0.05
par    c4                   nui      0.004        0.001      -0.05    0.05
par    c5                   nui      0.004        0.001      -0.05    0.05
par    Aps100x100           nui      2.5E-04      1E-05      0.0      0.1
par    Aps100x143           nui      1.1E-04      7E-06      0.0      0.1
par    Aps100x217           nui      9.9E-05      6E-06      0.0      0.1
par    Aps143x143           nui      4.7E-05      2E-06      0.0      0.1
par    Aps143x217           nui      3.1E-05      3E-06      0.0      0.1
par    Aps217x217           nui      7.6E-05      6E-06      0.0      0.1
par    Asz                  nui      1            0.1        0.0      10
par    Acib                 nui      1.           0.1        0.0      10
par    AdustTT              nui      1            0.1        0.0      2
fix    AdustPP              nui      0.00         0.1        0.0      2
fix    AdustTP              nui      0.00         0.1        0.0      2
```

```
par    Aksz                    nui      0.00          1.0          0.0      10
par    Aszxcib                 nui      0.00          1.0          0.0      10

#PRIORS
gauss1      AdustTT         1         0.2
gauss1      Acib            1         0.2
gauss1      A_planck        1         0.0025
gauss1      c0              0         2e-3
gauss1      c1              0         2e-3
gauss1      c3              0         2e-3
gauss1      c4              0.002     2e-3
gauss1      c5              0.002     2e-3
#------------------------------


#------------------------------
#CMB Very-high-l (SPT)
SPT_High=HighEll/SPT_high_2014.lik
par       SPT_high_95_cal           nui       0.9961   0.002    0.9       1.1
par       SPT_high_150_cal          nui       1.002    0.002    0.9       1.1
par       SPT_high_220_cal          nui       1.015    0.002    0.9       1.1
par       SPT_high_Aps_95x95        nui       7.425    0.01     0.1       50
par       SPT_high_Aps_95x150       nui       5.147    0.01     0.1       50
par       SPT_high_Aps_95x220       nui       8.8      0.01     0.1       50
par       SPT_high_Aps_150x150      nui       6.649    0.01     0.2       50
par       SPT_high_Aps_150x220      nui       14.15    0.01     1.5       50
par       SPT_high_Aps_220x220      nui       36.07    0.01     3         200
fix       SPT_ADust                 nui       1

#PRIORS
gauss1 SPT_high_95_cal  1.01  .01
gauss1 SPT_high_150_cal  1.01  .01
gauss1 SPT_high_220_cal  1.01  .02
#------------------------------


#------------------------------
#CMB Very-high-l (SPT-low)
SPT_Low     = HighEll/SPT_low.lik
par                 SPT_low_Aps               nui               20.32             0.01              1.
                      60
par                 SPT_low_cal               nui               1.00              0.01              0
                        2
#------------------------------


#------------------------------
#CMB Very-high-l (ACT)
ACT_equat=HighEll/ACT_equat.lik
par       ACT_equat_148_cal         nui       0.9991   0.002    0.9       1.1
par       ACT_equat_220_cal         nui       1.013    0.002    0.9       1.1
par       ACT_equat_ADust           nui       1.719    0.01     0.05      10
par       ACT_equat_Aps_148x148     nui       7.159    0.01     0.1       50
par       ACT_equat_Aps_148x220     nui       20       0.01     0.1       50
par       ACT_equat_Aps_220x220     nui       60       0.01     10        150
ACT_south=HighEll/ACT_south.lik
par       ACT_south_148_cal         nui       1.007    0.002    0.9       1.1
par       ACT_south_220_cal         nui       1.032    0.002    0.9       1.1
par       ACT_south_ADust           nui       1.3      0.01     0.05      10
par       ACT_south_Aps_148x148     nui       9        0.01     0.1       50
par       ACT_south_Aps_148x220     nui       16.29    0.01     0.1       50
par       ACT_south_Aps_220x220     nui       60       0.01     10        150
#------------------------------
```

```
#------------------------------
#CMB lensing (Planck)
lensing_file=planck_data/lensing/smica_g30_ftl_full_pttptt.clik_lensing#
#------------------------------


#------------------------------
#SNIa (JLA)
JLA_SNIA_File=JLA/JLA.lik
par         Mabs            nui     -19.04          .03             -19.25          -18.85
par         alpha           nui     0.141           .01             0.10            .18
par         beta            nui     3.101           0.1             2.5             3.7
par         DeltaM          nui     -.076           0.03            -.13            -0.01
#------------------------------


#------------------------------
#BAO
BAOFile=BAO/BAO1D_new.lik
BAO2DFile=BAO/Boss_anderson_dr11.lik
#------------------------------


###################################################################




###################################################################
#fitter
nitermax=50000
set_stra=2
set_tol=0.00001

#output options
remove_cosmo_limits=false
doHesse=true
###################################################################
```